

19
8-103

UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:

STEPHEN FULD

Serial No.: 09/479,146

Filed: January 7, 2000

Group Art Unit: 2184

Examiner: Michael Maskulinski

For: METHOD AND SYSTEM FOR RECONSTRUCTING
DATA SERIALY ARRANGED ON A MAGNETIC TAPE TRACK (As Amended)

Attorney Docket No.: 99-051-TAP (STK 99051 PUS)

RECEIVED

APPEAL BRIEF

AUG 05 2003

Technology Center 2100

Mail Stop Appeal Brief - Patents
Commissioner for Patents
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an appeal brief (submitted in triplicate) in support of an appeal from the final rejection of claims 12-14 and 16-18 in the final Office Action mailed on May 30, 2003.

08/05/2003 DTESSEM1 00000004 09479146

01 FC:1402

320.00 OP

I. REAL PARTY IN INTEREST

The real party in interest is Storage Technology Corporation ("the Assignee"), a corporation organized and existing under the laws of the state of Delaware, and having a place of business at One StorageTek Drive, MS-4309, Louisville, Colorado 80028-4309. The Assignee's interest in this patent application is set forth in an assignment recorded in the U.S. Patent and Trademark Office on January 7, 2000 at Reel 010489, Frame 0615.

CERTIFICATE OF MAILING UNDER 37 C.F.R. § 1.8

I hereby certify that this paper, including all enclosures referred to herein, is being deposited with the United States Postal Service as first-class mail, postage pre-paid, in an envelope addressed to: Mail Stop Appeal Brief - Patents, Commissioner for Patents, U.S. Patent & Trademark Office, P.O. Box 1450, Alexandria, VA 22313-1450 on:

August 1, 2003
Date of Deposit

James N. Kallis
Name of Person Signing

Signature

II. RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to the Applicant, the Applicant's legal representative, or the Assignee which will directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 12-14 and 16-18 are pending in this application. Claims 12-14 and 16-18 (reproduced for reference in the attached Appendix) have been finally rejected and are the subject of this appeal. Of claims 12-14 and 16-18, claims 12 and 16 are independent claims.

IV. STATUS OF AMENDMENTS

The Applicant filed an Amendment after Final on June 26, 2003 in which the Applicant amended the title. The Advisory Action mailed on July 16, 2003 in response to the Amendment after Final did not indicate whether the amended title has been entered. The Applicant has identified this patent application in the appeal papers with the title as amended.

V. SUMMARY OF THE INVENTION

The claimed invention is generally directed to a method and system for reading data blocks from magnetic tape in which the data blocks and a parity block are serially arranged on a track (14) of the magnetic tape. The claimed invention accumulates a running parity of the data blocks as the data blocks are read sequentially from the magnetic tape track. By accumulating a running parity, good data blocks read prior to a bad data block do not have to be reread for obtaining the parity information of the good blocks which is required to

complete the reconstruction of the bad data block. (See page 1, lines 4-9; page 4, lines 1-17; page 10, line 3 through page 11, line 2; and FIGS. 2-4 of the Applicant's specification.)

The claimed invention employs Redundant Array of Independent Disks (Tapes) (RAID, RAIT) principles for reconstructing a bad data block from the remaining data blocks and a corresponding parity block in a given data stripe (20). (See FIG. 2 and page 9, line 3 through page 10, line 2 of the Applicant's specification.) Particularly, the claimed invention employs RAID-4 principles. Under the RAID-4 standard, during writing operations the data blocks of a given data stripe having N data blocks are striped across N disk/tape data drives and then a parity block is striped to a disk/tape parity drive. The parity block represents the parity of the N data blocks. This process is repeated for each given data stripe (22). During reading operations, the data drives and the parity drive are read sequentially to read the data blocks and parity block in order. If all of the data drives are operable then the data blocks may be read directly from the data drives. If one of the data drives fails then the parity information of the parity drive may be used in conjunction with the remaining data blocks of the other data drives in order to reconstruct the data block in the failed data drive.

The claimed invention transfers the RAID-4 standard to a single track of a magnetic tape, i.e., "single tape RAIT". (See page 8, lines 3-18 of the Applicant's specification.) Accordingly, N data blocks and a corresponding parity block of a given data stripe (20) are serially arranged on a track (14) of the magnetic tape. This arrangement may then be repeated with N data blocks and a corresponding parity block of the next given data stripe (22) serially arranged on the magnetic tape track after the given data stripe (20) and so on.

With reference to FIG. 4 of the Applicant's drawings and page 10, line 3 through page 12, line 22 of the Applicant's specification, the data blocks are read sequentially from the magnetic tape track as indicated by block 42. The data block currently being read

is determined to be good or bad as indicated by block 44. The parity of the good data blocks are accumulated as the data blocks are being read as indicated by block 46. The data block currently being read is sent to the host if the currently being read data block does not follow a bad data block as indicated by block 47. The parity block is read after all of the data blocks in the given data stripe have been read as indicated by block 48. If one of the data blocks is bad, the bad data block is reconstructed from the accumulated parity of the data blocks and the parity block in order to form a reconstructed good data block as indicated by block 50. The reconstructed good data block is then sent to the host as indicated by block 52. The good data blocks after the reconstructed good data block are sent to the host in sequential order as indicated by block 54.

The claimed invention transfers the RAID-4 standard to a single track of a magnetic tape because a common failure mode when reading magnetic tape is the loss of a block on a magnetic tape track. By providing the parity block based on the data blocks of a given data stripe in the manner of the RAID-4 standard, the claimed invention employs a well-known solution for reconstructing a bad data block in a single magnetic tape environment.

In summary, the implementation of the "single tape RAIT" in accordance with the claimed invention is intended to represent the idea of data blocks and a corresponding parity block of a given data stripe serially arranged on a tape track. This is analogous to the typical RAID/RAIT environments in which each data block of a serially arranged set of data blocks is striped to a respective disk/tape data drive with the corresponding parity block also being striped to a respective disk/tape parity drive.

VI. ISSUE

The Examiner rejected claims 12-14 and 16-18 under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,018,778 issued to Stolowitz ("Stolowitz") in view

of White, How Computers Work ("White"). The sole issue on appeal is whether Stolowitz in view of White makes a *prima facie* showing of obviousness of claims 12-14 and 16-18.

VII. GROUPING OF CLAIMS

Claims 12-14 and 16-18 stand or fall together.

VIII. ARGUMENT

To establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference must teach or suggest all the claim limitations. MPEP 2143. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, not in applicant's disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991).

With respect to the last requirement, to establish *prima facie* obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974). "All words in a claim must be considered in judging the patentability of that claim against the prior art." *In re Wilson*, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970). If an independent claim is nonobvious under 35 U.S.C. § 103, then any claim depending therefrom is nonobvious. *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988).

1. The Claimed Invention

The claimed invention, as recited in independent claims 12 and 16, is a method and an associated system for providing data blocks from a magnetic tape to a host. The method and system are for use in a "single magnetic tape RAIT" environment. Such a single magnetic tape RAIT environment includes a magnetic tape having data blocks and a parity block serially arranged on a single track of the magnetic tape with the parity block following the data blocks. The parity block is based on the data blocks as conventionally known.

The method includes reading the data blocks sequentially from the track of the magnetic tape and determining if the data block currently being read is good or bad based on the reading of the data block currently being read. The data block currently being read is provided to the host if the currently being read data block does not follow a bad data block. If one of the data blocks is bad, the method includes storing the good data blocks following the bad data block in sequential order.

Parity of the good data blocks is accumulated as the data blocks are being read. The parity block is then read from the track of the magnetic tape after all of the data blocks have been read. If one of the data blocks is bad, the bad data block is then reconstructed from the accumulated parity of the data blocks and the parity block in order to form a reconstructed good data block. The reconstructed good data block is then provided to the host and then the stored good data blocks are provided to the host in sequential order.

2. Stolowitz and White

The Examiner posited that Stolowitz discloses the claimed invention with the exception of explicitly disclosing a magnetic tape having data blocks and a parity block in which the data blocks and the parity block are serially arranged on the magnetic tape with the

parity block following the data blocks and being based on the data blocks. The Examiner posited that White discloses the format of a tape having parallel tracks with each track being divided into segments of blocks of bytes. Specifically, each track being divided into 512 or 1,024 bytes and each segment typically containing 32 blocks. The Examiner posited that eight blocks in a segment contain error-correction codes and, as a result, the tracks comprise both data and parity.

3. The Claimed Invention Compared to Stolowitz and White

The claimed invention generally differs from any combination of Stolowitz and White in that the claimed invention is directed to an implementation of RAIT on a track of a single magnetic tape. As such, the data blocks and the parity block are serially arranged on a track of the magnetic tape with the parity block following the data blocks. The data blocks are read sequentially from the track of the magnetic tape and then the parity block is read from the magnetic tape after all of the data blocks have been read.

The claimed invention uses a parity block which is redundancy information calculated from the data blocks and serially arranges the data blocks and the parity block on the magnetic tape track. The principle of the parity block is as follows: take "N" data blocks and from them compute the parity block so that there are now "N+1" blocks. If any one of the "N+1" blocks is unreadable, the unreadable block can be reconstructed from the remaining "N" blocks, regardless of which block is unreadable. The calculation for determining the parity block and reconstructing an unreadable block is typically the logical "exclusive OR" or "XOR" operation as recited in dependent claims 13-14 and 17-18. Accordingly, the claimed invention transfers the RAID-4 standard to a single track of a magnetic tape. As indicated above, by providing the parity block based on the data blocks of a given data stripe in the manner of the RAID-4 standard, the claimed invention employs a well-known solution for reconstructing a bad data block in a single magnetic tape environment.

Again, the claimed invention is analogous to the typical RAID/RAIT environments in which each data block of a serially arranged set of data blocks is striped to a respective disk/tape data drive with the corresponding parity block also being striped to a respective disk/tape parity drive.

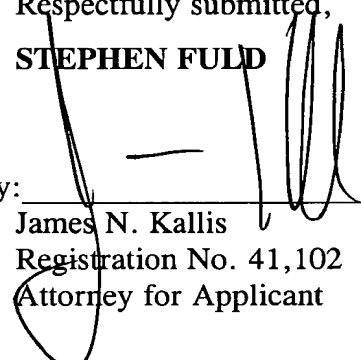
In contrast, White teaches the use of an error-correction code (ECC) in some of the blocks of a segment in a track and teaches the use of cyclic redundancy codes (CRC) for each block. In each case, the ECC and CRC are part of or appended to a data block. For example, White describes in point "2" how EC codes are computed and then appended to the end of the data and that after this data is transferred the process is repeated with "the next block of data from the disk." White further describes in the section entitled "Restoring Files" that the "controller computes a CRC code for each block and compares it with the CRC code written at the end of the block. If there's a discrepancy, error-correction routines usually can fix the data using the EC codes appended to each data block." White describes in point "5" that some blocks in a segment "contain error-correction codes" and that additionally "at the end of each block, the drive computer a cyclic redundancy check (CRC) for further error correction and appends it to the block."

Accordingly, White does not teach or suggest data blocks and a parity block serially arranged on a track of the magnetic tape with the parity block following the data blocks in the manner of the RAID-4 standard as claimed. In contrast, White teaches blocks serially arranged on a magnetic tape track with some of the blocks being data blocks and some of the blocks being data and ECC blocks. As such, if a block containing both data and the ECC were unreadable, then it is not clear as to how White would be able to construct the data in this unreadable block. In contrast, the claimed invention enables any unreadable block to be constructed from all of the remaining blocks. Therefore, modifying Stolowitz with White would not result in the claimed invention.

IX. SUMMARY

For these reasons, the Applicant respectfully submits that claims 12-14 and 16-18 are patentable under 35 U.S.C. § 103(a) over Stolowitz and White.

Respectfully submitted,
STEPHEN FULD

By: 
James N. Kallis
Registration No. 41,102
Attorney for Applicant

Date: August 1, 2003

BROOKS KUSHMAN P.C.
1000 Town Center, 22nd Floor
Southfield, MI 48075
Phone: 248-358-4400
Fax: 248-358-3351

Enclosure - Appendix

APPENDIX - CLAIMS ON APPEAL

12. In a magnetic tape having data blocks and a parity block in which the data blocks and the parity block are serially arranged on a track of the magnetic tape with the parity block following the data blocks and the parity block being based on the data blocks, a method for providing the data blocks from the track of the magnetic tape to a host, the method comprising:

reading the data blocks sequentially from the track of the magnetic tape;

determining if the data block currently being read is good or bad based on the reading of the data block currently being read;

providing the data block currently being read to the host if the currently being read data block does not follow a bad data block;

if one of the data blocks is bad, storing the good data blocks following the bad data block in sequential order;

accumulating parity of the good data blocks as the data blocks are being read;

reading the parity block from the track of the magnetic tape after all of the data blocks have been read;

if one of the data blocks is bad, reconstructing the bad data block from the accumulated parity of the data blocks and the parity block in order to form a reconstructed good data block;

providing the reconstructed good data block to the host; and

providing the stored good data blocks to the host in sequential order after the reconstructed good data block has been provided to the host.

13. The method of claim 12 wherein:

accumulating parity of the good data blocks includes exclusive ORing the parity of the good data blocks read prior to the good data block currently being read with the good data block currently being read.

14. The method of claim 13 wherein:

reconstructing the bad data block includes exclusive ORing the parity of the good data blocks with the parity block.

16. A data storage array system for providing data blocks to a host, the system comprising:

magnetic tape having data blocks and a parity block in which the data blocks and the parity block are serially arranged on a track of the magnetic tape with the parity block following the data blocks and the parity block being based on the data blocks;

a controller for reading the data blocks sequentially from the track of the magnetic tape and for reading the parity block from the track of the magnetic tape, wherein the controller determines if the data block currently being read is good or bad based on the reading of the data block currently being read, the controller providing the data block currently

being read to the host if the currently being read data block does not follow a bad data block, the controller reading the parity block from the track of the magnetic tape after all of the data blocks have been read;

a buffer, wherein if one of the data blocks is bad, the buffer stores the good data blocks following the bad data block in sequential order; and

a parity accumulator for accumulating parity of the good data blocks as the controller reads the data blocks;

wherein if one of the data blocks is bad, the controller reconstructs the bad data block from the accumulated parity of the good data blocks and the parity block in order to form a reconstructed good data block;

wherein the controller provides the reconstructed good data block to the host and then provides the good data blocks stored in the buffer to the host in sequential order after the reconstructed good data block has been provided to the host.

17. The system of claim 16 wherein:

the parity accumulator accumulates parity of the good data blocks by exclusive ORing the parity of the good data blocks read prior to the good data block currently being read with the good data block currently being read.

18. The system of claim 17 wherein:

the controller reconstructs the bad data block by exclusive ORing the parity of the good data blocks with the parity block.